Snapshot, Clone, and Boot from volume

Kazutaka Morita: NTT Cyber Space Labs.
<morita.kazutaka@lab.ntt.co.jp>
Isaku Yamahata: VA Linux Systems Japan K.K
<yamahata@valinux.co.jp>
<yamahata@private.email.ne.jp>

Openstack design summit: April 28, 2011

Snapshot and Clone

- Snapshot
  - Create a self-consistent image of the volume
  - You can use snapshots to take backups

```
$ ec2-create-snapshot volume_id
```

- Clone
  - Create a new volume from the snapshot
  - You can create lots of same volumes from one golden snapshot

```
$ ec2-create-volume –snapshot snapshot_id -z zone
```

Implementation

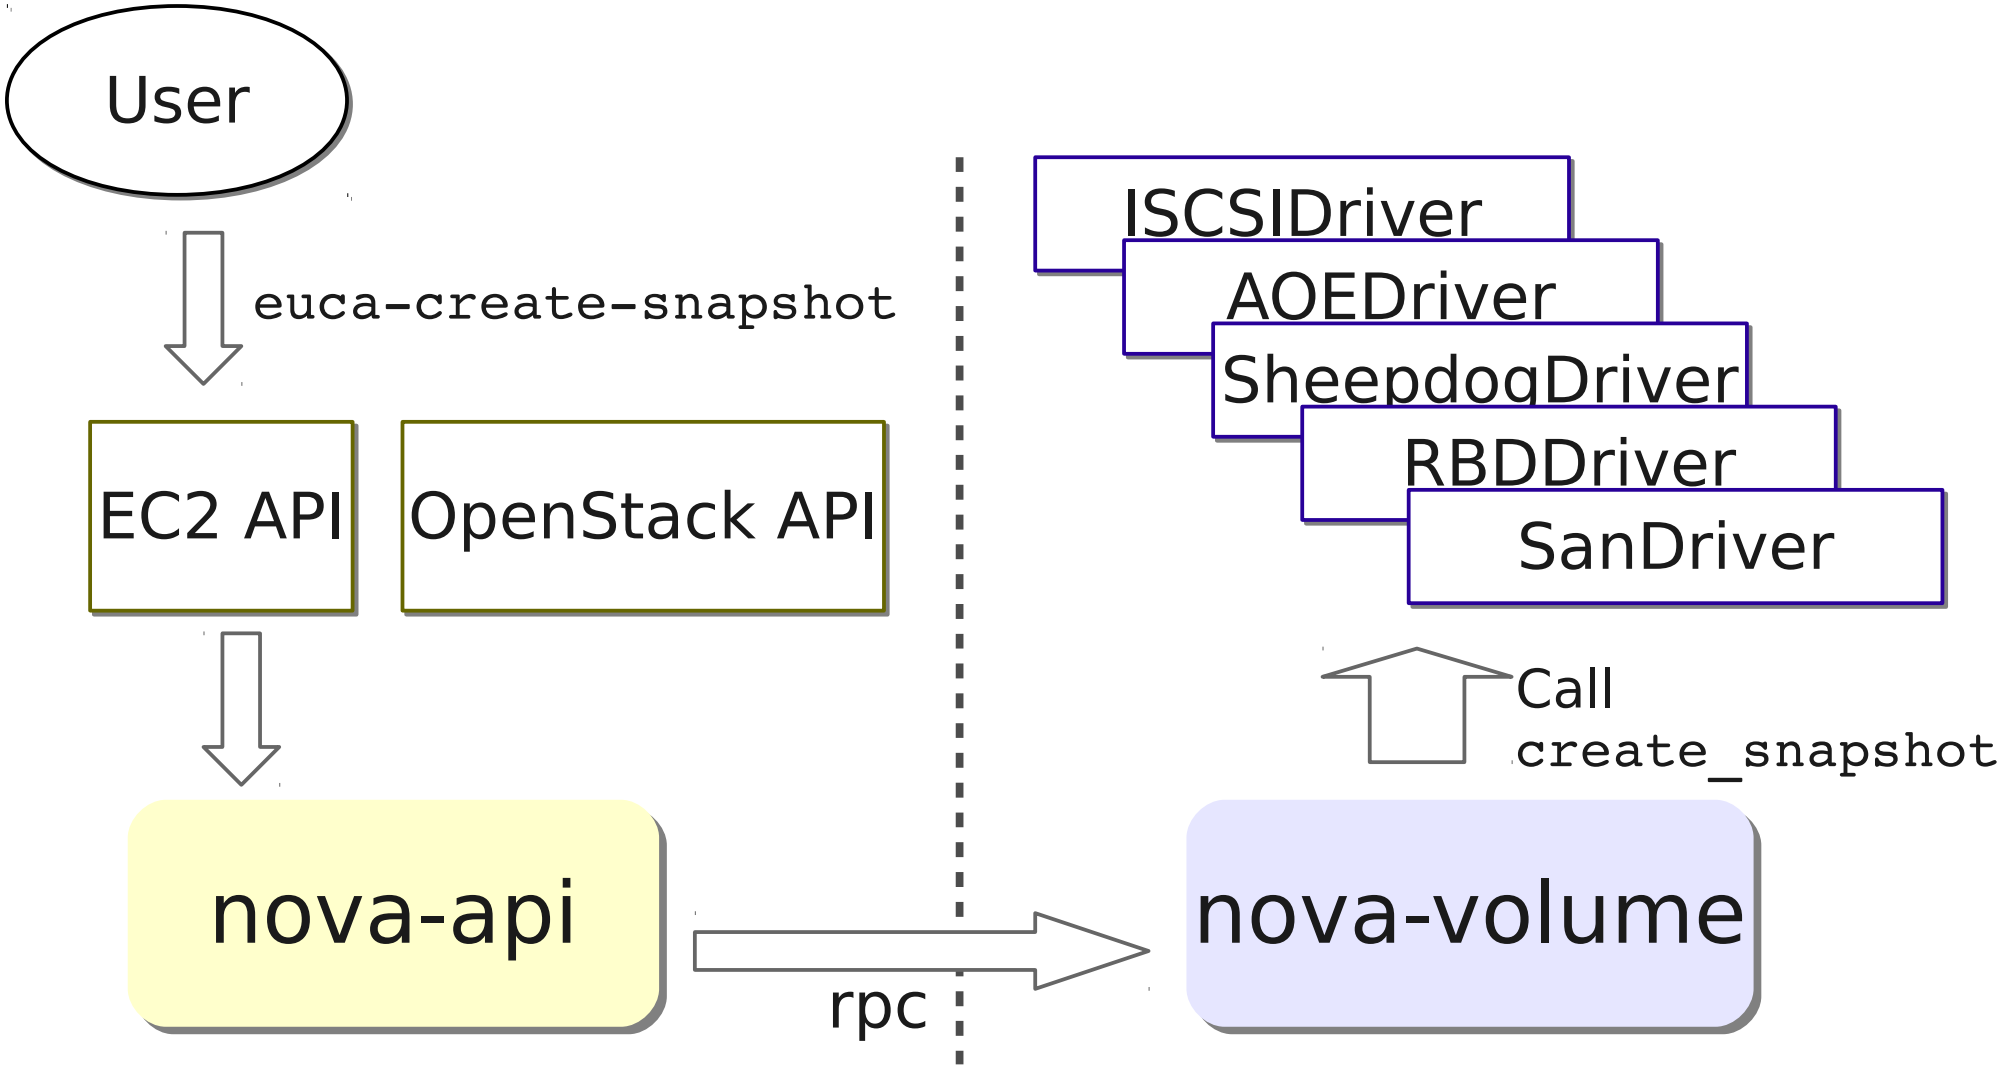

New methods to volume driver

```
class VolumeDriver(object):

    ...

    def create_snapshot(self, snapshot):
        """Creates a snapshot."""
        raise NotImplementedError()

    def delete_snapshot(self, snapshot):
        """Deletes a snapshot."""
        raise NotImplementedError()

    def create_volume_from_snapshot(self, volume, snapshot):
        """Creates a volume from a snapshot."""
        raise NotImplementedError()
```

DB schema changes (1)

- Create new table *Snapshot*

```
snapshots = Table('snapshots', meta,
        Column('created_at', DateTime()),
        Column('updated_at', DateTime()),
        Column('deleted_at', DateTime()),
        Column('deleted', Boolean()),
        Column('id', Integer(),  primary_key=True),
        Column('volume_id', Integer()),
        Column('user_id', String()),
        Column('project_id', String()),
        Column('status', String()),
        Column('progress', String()),
        Column('volume_size', Integer()),
        Column('scheduled_at', DateTime()),
        Column('display_name', String()),
        Column('display_description', String())
        )
```

DB schema changes (2)

- adds a new column *snapshot_id* to the table *volumes*

```
snapshot_id = Column('snapshot_id', Integer())
```

Where to store snapshots

- Amazon EC2 stores EBS snapshots to S3
- Where should Nova store volume snapshots?
  - To Swift?
  - To Glance?
  - To nothing? (only stored in nova-volume)

Snapshot/clone with Linux LVM2

- Linux LVM2 doesn't support snapshot of snapshot.
- If deleting the origin volume, the dependent snapshots will also be deleted
  - EC2 allows delete volume/snapshot independently: abandon this feature for now

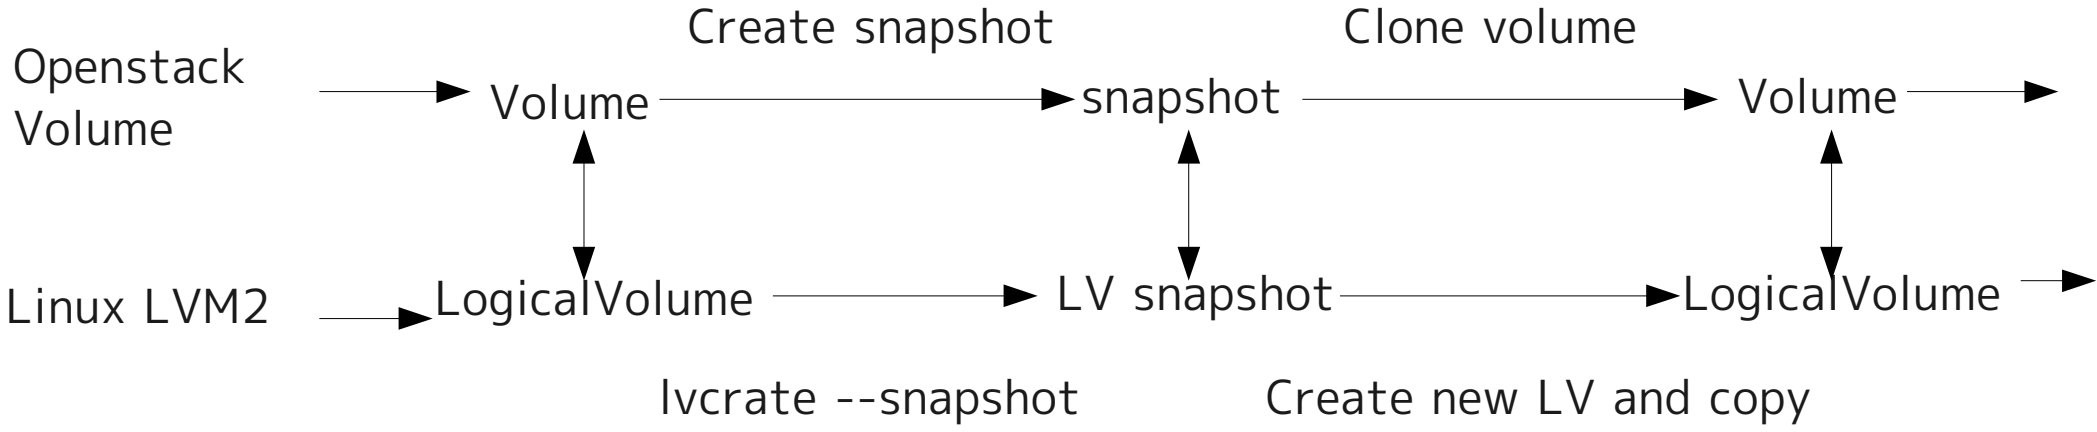

Boot from volume(EBS boot)

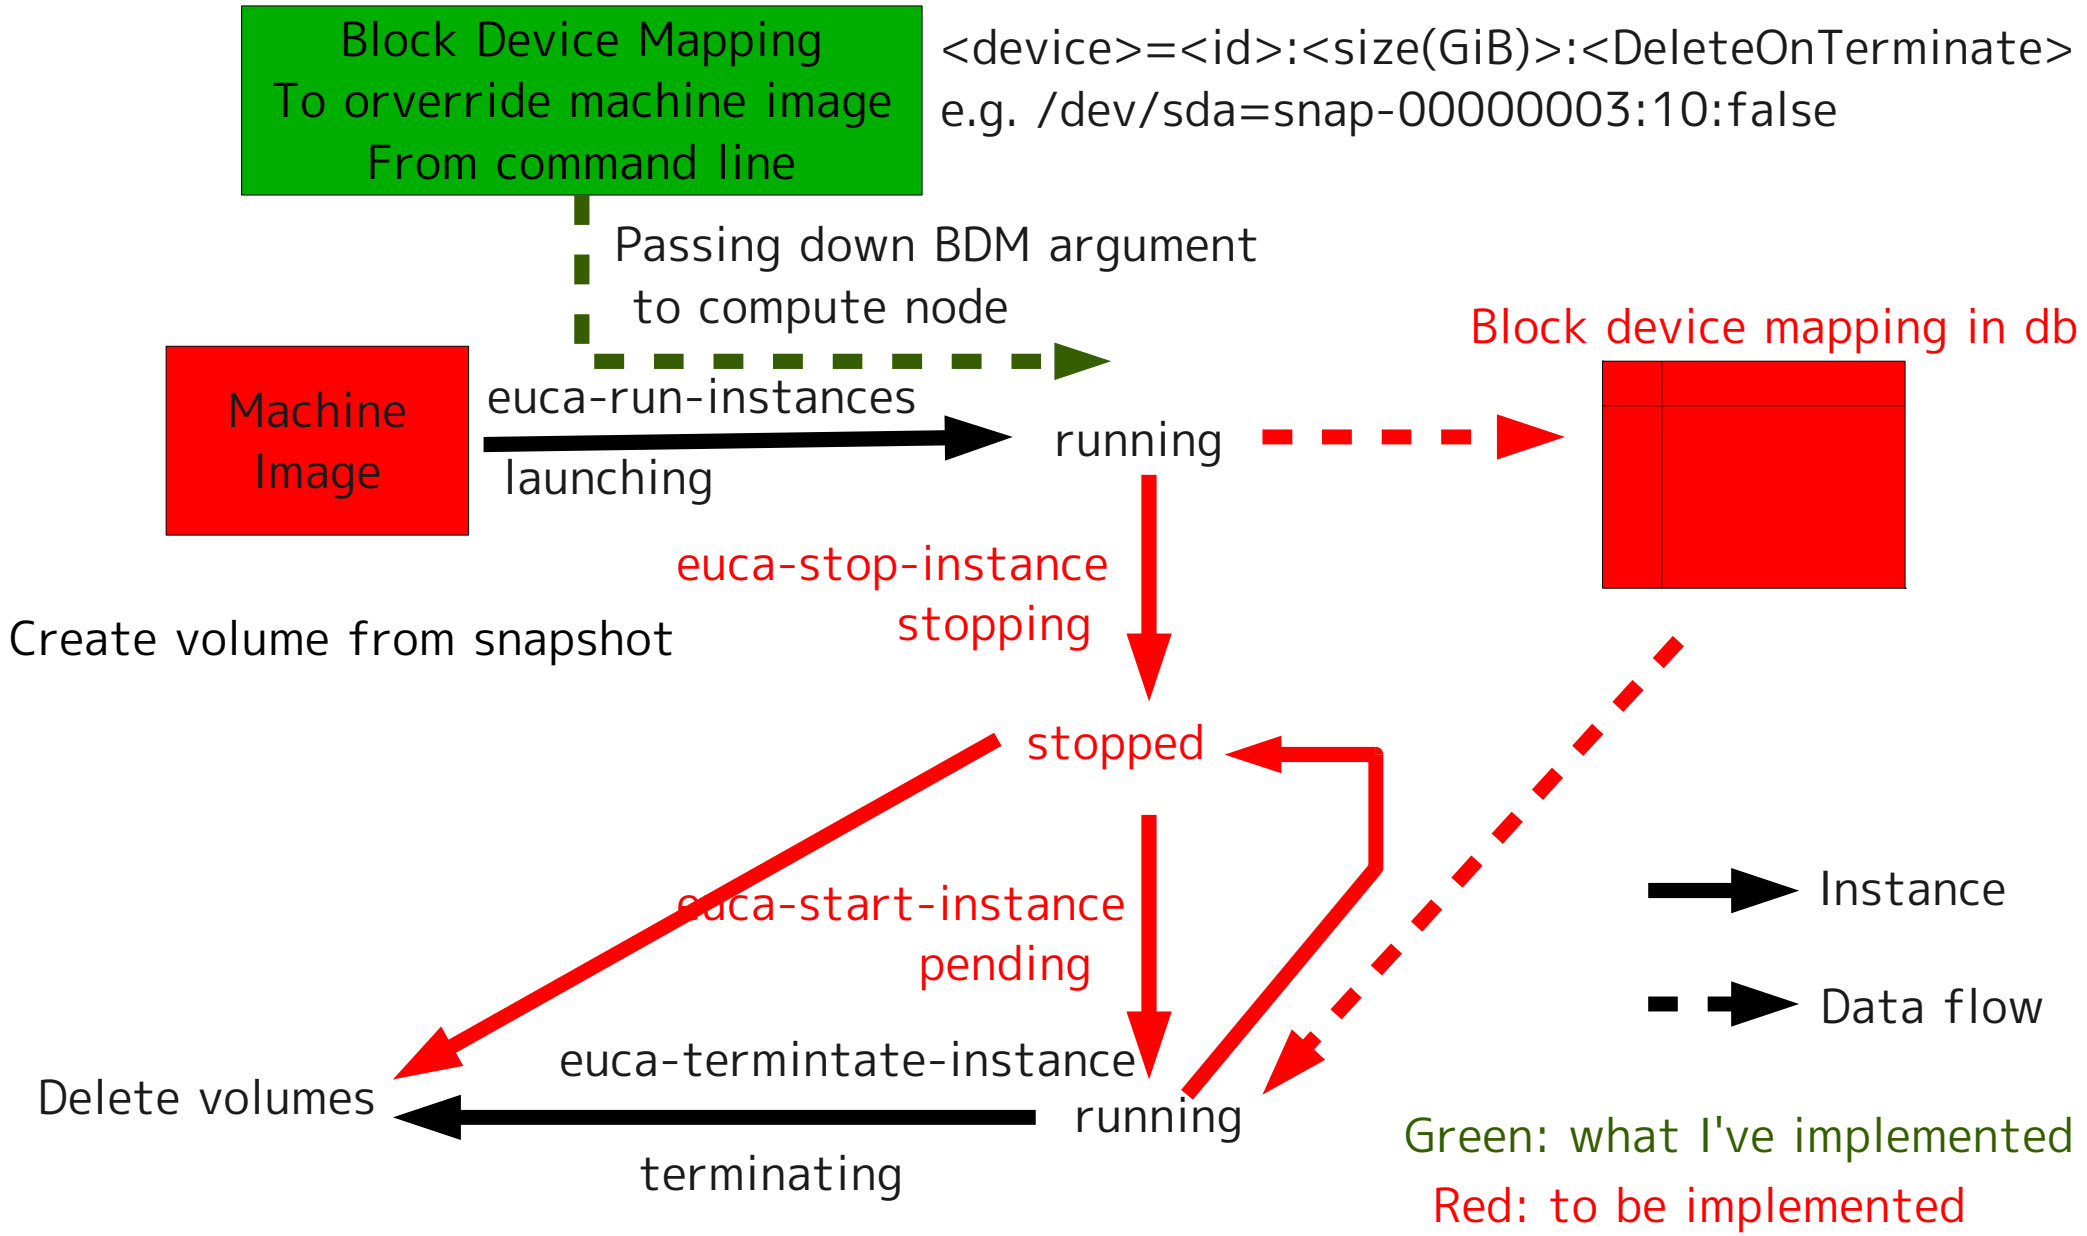

Machine Image changes

- +Machine Image Type: S3 or EBS
- +Block device mapping
  - +Snapshot id or volume id or virtual device
    - EC2 allows only snapshot id
  - +Volume Size
  - +DeteleteOnTerminate
- +Root Device
- ec2-create-image
- ec2-register

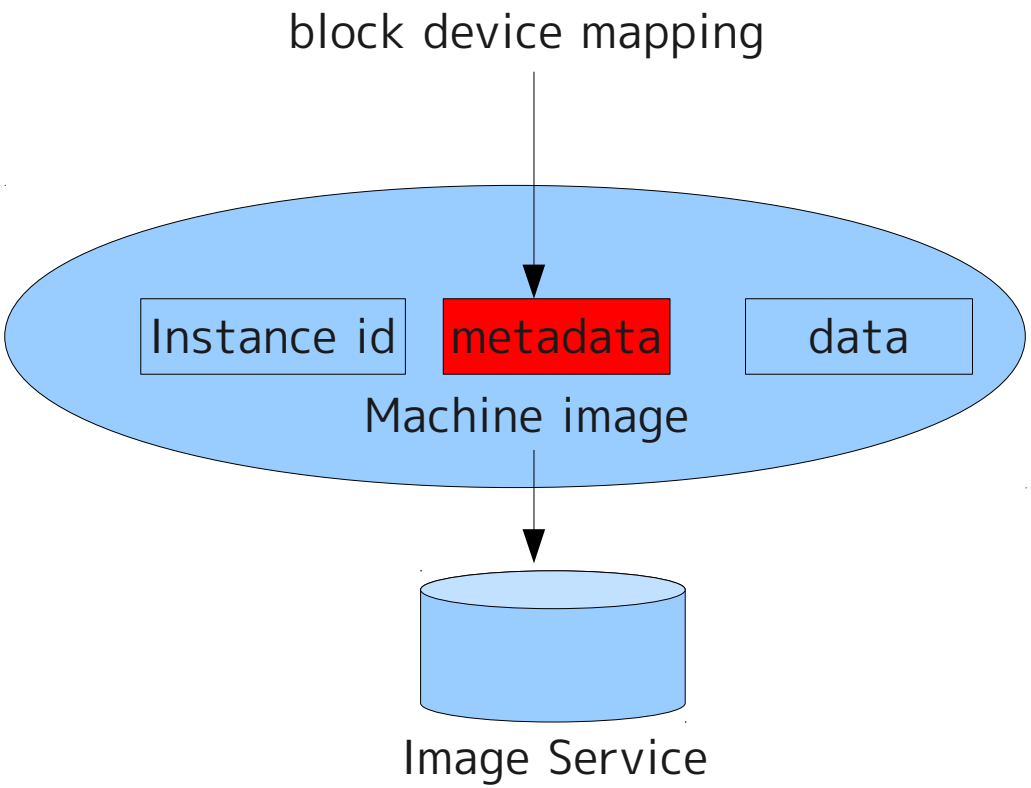

DB schema changes

- Volume Table
  - +DeleteOnTerminate
- +BlockDeviceMapping Table
  - +instance id
  - +Device Name(device name in guest)
  - +Mapped device
    - +Virtual device(ephemeral0...) local storage which isn't persistent
    - +Volume
      - Snapshot id or Volume id
      - Volume Size in GiB
      - DeleteOnTerminate: bool
    - +No Device

start/stop Instance

- start/stop instance
  - Mostly same to run/terminate instance except db manipulation
- New status
  - Stopping/stopped/pending